End-to-end science operations in the era of Extremely Large Telescopes

Olivier R. Hainaut,[a,]* Marie Lemoine-Busserolle,[b] Christophe Dumas,[c] Robert W. Goodrich,[d] Bryan W. Miller,[b] Michael F. Sterzik,[a] Thomas Bierwirth,[a] Sidney Wolff,[b] Andrew W. Stephens,[b] Gelys Trancho,[c] Warren Skidmore,[c] and Kim Gillies[c]

aEuropean Southern Observatory, Garching bei München, Germany
bNSF's NOIRLab, Tucson, Arizona, United States
cTMT International Observatory, Pasadena, California, United States
dGMTO Corporation, Pasadena, California, United States

Abstract. Observatory end-to-end science operations is the overall process starting with a scientific question, represented by a proposal requesting observing time; ending with the analysis of observation data addressing that question; and including all of the intermediate steps needed to plan, schedule, obtain, and process these observations. Increasingly complex observing facilities demand a highly efficient science operations approach and at the same time must be user friendly to the astronomical user community and enable the highest possible scientific return. Therefore, this process is supported by a collection of tools. We describe the overall end-to-end process and its implementation for the three upcoming Extremely Large Telescopes (ELTs): European Southern Observatory's ELT, the Thirty Meter Telescope, and the Giant Magellan Telescope. © *The Authors. Published by SPIE under a Creative Commons Attribution 4.0 International License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI.* [DOI: 10.1117/1.JATIS.8.2.021508]

1 End-to-End Operations

1.1 *History*

Historically, astronomers obtained observing time on their local or national telescope, travelled there to perform the observations (sometimes involving an arduous trip and an extended stay), and returned to their institute with stacks of photographic plates. Over time, the main change to this paradigm was to replace plates with electronic detectors that recorded their output on magnetic tapes or optical disks. Each observer would acquire the ancillary data required to calibrate their observations and, although the plates (or tapes) formally remained the property of the observatory, little or no effort was made to make these available to the broader community. Some exceptions existed, e.g., the extensive Harvard plate collection.[1]

For ground-based astronomy, this changed with the 8- to 10-m telescopes built in the 1990s.[2] In addition to the traditional visitor mode (VM, also known as classical observing) described above, European Southern Observatory (ESO), Gemini, and others introduced service mode (SM, also known as queue mode): observations are defined in advance, then performed by observatory staff who optimized the schedule of the telescope and matched the observing conditions with the requirement of each observation. In parallel, the observatory started to calibrate the instruments systematically and consistently (setting up detailed calibration plans for all instruments) and to methodically archive all data produced.

The scientific process of observational astronomy is cyclical, and it was realized that the new observing modes would only be successful if observing program information was managed from start to end. Quinn[3] and Puxley et al.[4] described a similar process that starts with a scientific

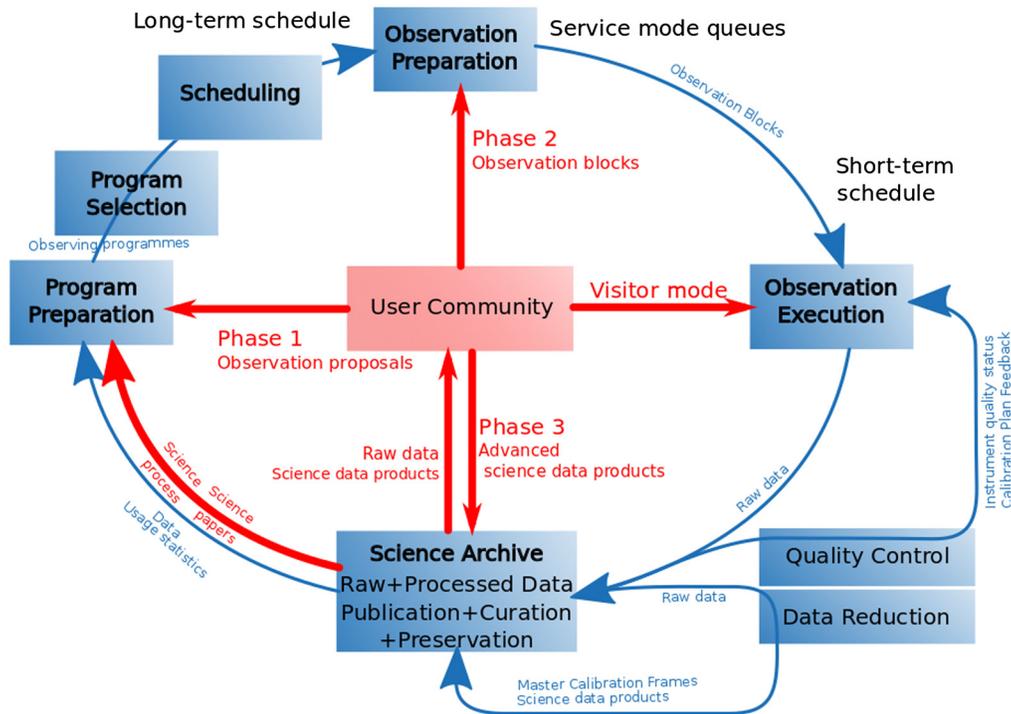*Address all correspondence to Olivier R. Hainaut, ohainaut@eso.org

Fig. 1 End-to-end operation process, as implemented at ESO (and very similarly for the other ELTs). The blue boxes represent the main sub-processes; the arrows are labeled with the information flowing through them. The red arrows show the main interactions with the community. Evolved from Ref. 3.

question addressed by a request for observing time; proceeds with the observations being prepared, executed, processed, analyzed; and finally, answers the scientific questions (hopefully) with a publication, which triggers new questions. This end-to-end process enables efficient science operations of complex instrumentation and at the same time enables an optimal scientific return for the community. ESO's current implementation is shown in Fig. 1. The concepts that it describes are relevant to most observatories.

This paper reviews and compares how the end-to-end operations are foreseen for the next generation of Extremely Large Telescopes (ELTs).

1.2 *Evolution*

At ESO, the end-to-end process was formalized for the Very Large Telescope (VLT) in the 1990s as the original data flow system (DFS) described by Quinn[3] and deployed on the New Technology Telescope (NTT) in 1996. It was subsequently upgraded to become an operational prototype of the VLT. The DFS was—and still is—built around the concept of the observation block (OB), which is the atom holding all of the information required to perform an independent set of observations on a target. An OB consists of a series of templates, each defining a type of observation [e.g., to acquire a set of adaptive optics (AO) images with offsets or to acquire long-slit spectra on and off-target]. The OB only lists the template parameters (e.g., exposure times, filter, grating and wavelength, size and number of offsets, AO configuration, etc.); the actual sequence of instruction to be performed using these parameters exists only in the instrument. The OB is the atomic unit defining the minimum dataset making scientific sense, and the template is the indivisible quantum in terms of the calibration plan (as an example, the OB includes templates describing a series of images in R and V, for a program aiming at measuring colors; for the calibration plan, the R and the V filters must both be calibrated). The templates, the definition of their parameters, and their allowed ranges (together with a series of additional configurations files, defining, for instance, the relevant atmospheric constraints), constitute the instrument package (IP), which fully describes the characteristics and capabilities of the instrument.

Originally, the various subprocesses shown in Fig. 1 were implemented with relatively independent tools. Over 25 years of operation, the DFS evolved, integrating new concepts (e.g., scheduling containers grouping OBs, for instance, for large surveys on the VST and VISTA), new observing modes (e.g., rapid response mode for targets of opportunity requiring extremely urgent observations), new instruments with special requirements (e.g., Echelle SPectrograph for Rocky Exoplanets and Stable Spectroscopic Observations (ESPRESSO), which can be used during the same night on any of the four VLT units or can use all four simultaneously), and the production of science-ready processed data. The tools grew organically, and some of the underlying technologies aged or even became obsolete. Nevertheless, the DFS still provides a common interface to around 25 instruments mounted on the four VLTs and several other smaller telescopes. This homogeneous system in science operations is the counterpart to the control system; it is homogeneous over all systems in terms of hardware (via the VLT Standard) and software (due to the standard Central Control Software).

In preparation for the ELT, ESO reviewed in depth the concepts at the core of the DFS, considering the new requirements from the ELT and the stipulation that the ELT must be fully integrated with the VLT operations. Requirements from new instruments (e.g., AO and laser guide stars were at their infancy at the beginning of the VLT, whereas they are now ubiquitous), technology obsolescence (and the opportunities provided by recent technologies), and—not least—requests from the community were also considered. Highlights of this review are summarized in a series of messenger articles (Ref. 5 for SM; Ref. 6 for the archive; Ref. 7 for Phase 3; and Refs. 8 and 9 for scientific return). The result of the review was that the fundamental concepts of the DFS (in particular the IPs, OBs, and templates) were sound and would meet all of the new requirements. Although originally used only as an interface between the instrument control system and the observation handling system at the observatory, the use of IP is now being generalized across the entire DFS as the single, authoritative "hardware abstraction layer." It represents the physical properties of the instrument and reflects its actual (and eventually changing) capabilities, ensuring a consistency over the entire system that was lacking in the original implementation. This gives ESO the opportunity to modernize the tools implementing the DFS, integrating them better and building them to support operations over an expected lifetime of at least 20 years. A coherent development approach was considered with the onset of the "DMO/DFS program" in 2013 to coordinate these efforts; a high-level, if slightly outdated, description of this plan was published by Hainaut et al.[10]

The United States Extremely Large Telescope Program (US-ELTP) was formed as a collaboration of the TMT International Observatory, the Giant Magellan Telescope project (GMTO Corporation), and the National Science Foundation (NSF) National Optical-Infrared Astronomy Research Laboratory (NOIRLab). US-ELTP's main goal is to provide the user services to take full advantage of the capabilities of GMT and TMT by providing training, assistance, and integrated user-facing software tools. NSF's NOIRLab, together with TMT and GMT, developed a detailed plan for the user support system over all phases of the end-to-end process, called Scientific Data Life Cycle (SDLC) in this context and visualized in Fig. 2, and allocated the responsibility for each action that must take place. The Gemini Observatory is now a part of NSF's NOIRLab. Its expertise and the evolution of its operations therefore naturally served as a precursor on which the implementation of the SDLC tools builds extensively.

The evolution of the Gemini science operations followed a similar trajectory to that of ESO's. All of the original tools have evolved substantially, e.g., the Phase I Tool and International Time Allocation Committee software were rewritten, the archive was replaced with the Gemini Observatory Archive running on Amazon Web Service,[11] and the data reduction software is being transitioned from the Gemini IRAF package to the DRAGONS Python pipeline environment.[12] Gemini's original operations model envisioned 50% visitor/classical mode and 50% service/queue mode. However, investigators quickly voted for queue and now over 90% of the observing is done using some form of SM. Gemini then implemented new observing and proposal modes—rapid target-of-opportunity (ToO) mode with programmatic access, fast turnaround, and large-long programs—to enhance flexibility for programs of different sizes and to support student training.[13] Also, after years of usability improvements to the Observing Tool (OT), the point at which new required capabilities were too difficult to implement in the existing infrastructure
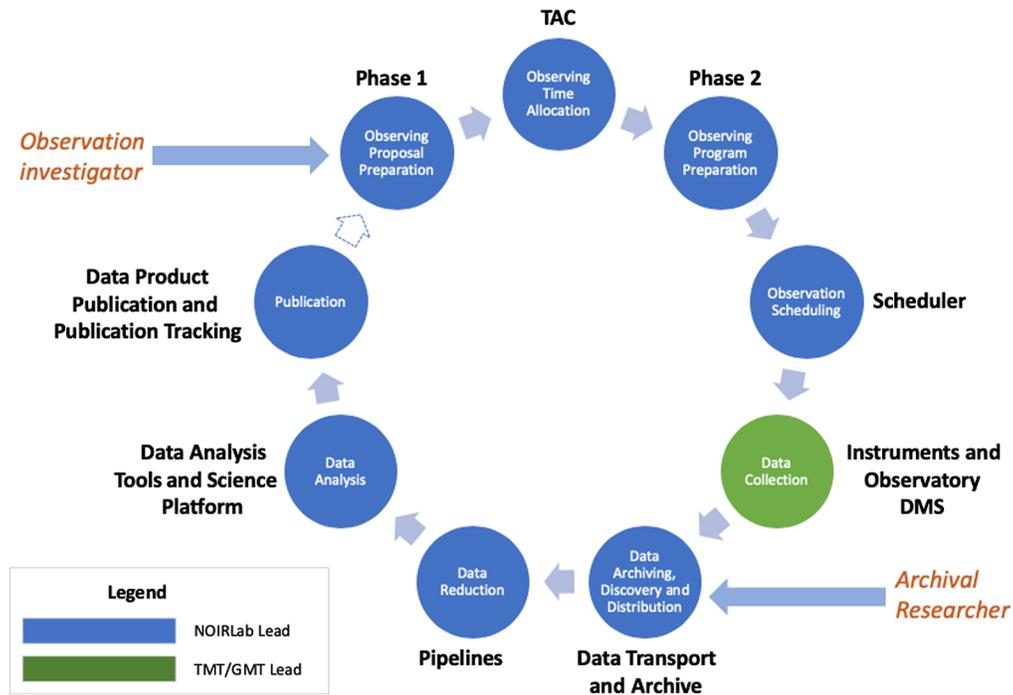
Fig. 2 The SDLC and the main phases that the UPP will support.

was reached. Therefore, NSF's NOIRLab/Gemini has initiated the Observatory Control System Upgrades Program to reimagine and replace most of the high-level operations applications with modern, web-based tools.[14] The Gemini Program Platform (GPP) is the core of this system and will unify proposal preparation, integration time calculators, and observation definition in a single web app called Explore. A single, cloud-based, observing database will unify Gemini North and South and allow them to work together via an automatic, real-time scheduler.

Building off the Gemini experience and work on GPP, NSF's NOIRLab will build the US ELT Program Platform (UPP), which will provide a significant subset of the functionality within the overall SDLC. The SDLC also includes, of course, the execution of observations and data collection by GMT and TMT. The UPP will provide a common user interface that will support applications for observing time on either or both the TMT and GMT and will make it possible to compare the performance of the imagers and spectrographs on each to determine the best strategy for specific programs. The development of these user services will be carried out by NSF's NOIRLab and TMT and GMT partners, as well as the US community. Similarities to NSF's NOIRLab/Gemini include access to two telescopes, one in each hemisphere; demanding observations that require special atmospheric conditions for AO and thermal infrared observations; a complex suite of instruments, several of which are available on any given night; and extensive use of queue and remote observing. Therefore, in the sections that follow, we will reference NSF's NOIRLab/Gemini current capabilities and planned improvements when we describe the end-to-end process for GMT and TMT.

In Sec. 2, we will go through the various sub-processes of the end-to-end science operations and compare the plans for ELT, GMT, and TMT.

1.3 Technologies

Over the past 20 years, tools supporting end-to-end operations have been implemented as software that users would install on their own computers. The wide zoo of operating systems used in the community resulted in complications for the developers (and often frustration for the users) and presented obstacles to the deployment of updates and bug fixes. New technologies, however, support platform-independent software access, via server-side software, exportable self-contained "containers," or other means.

ESO's strategic choice was to move toward online and web-based user interfaces that allow for efficient and platform-independent development. These interfaces are connected via application programming interfaces (APIs) to servers hosted by the organization. Scripts can also use the APIs, allowing power users to automate workflows requiring long or repetitive interactions with the web interfaces and even to develop custom user interfaces. Having the "business logic" running on ESO servers has the advantage that updates and bug fixes are deployed on the server and do not require the user to download or install new software. Also, with a strict separation between the user interface and the business logic, the interface can be updated or even completely overhauled without changing the server-side services. The technology powering web user interfaces evolves more rapidly than the core software technology.

NSF's NOIRLab's UPP will also include various web-based applications, web-based user interfaces, and a collection of secure APIs to all services to facilitate advanced user tasks, reporting, and automation. Some of these applications will be externally available to users whereas others will be for NSF's NOIRLab, GMT, and TMT internal use only. It is a requirement that most of the user-facing applications will be used for both GMT and TMT; therefore, a cloud-based deployment model has been selected.

1.4 *User Portal Infrastructure*

Many, if not most, of the interactions between the users and the services provided by the observatory benefit from the users identifying themselves, to have access to their personal information such as the proposals that they submitted or the data that they collected.

At ESO, this service is provided by the user portal,[15] which holds a profile for each user, including their basic identity, professional status, institution affiliation, year of PhD, and gender, as well as a series of keywords describing their scientific and technical expertise. The user's profile also includes the list of specific roles and privileges that they hold (for instance, referee for the TAC or operation staff astronomer), which are used to control access to the various confidential materials in the system. Keywords capture the scientific and technical expertise of each individual (used, e.g., to identify suitable referees for the TAC), and their affiliation to flag conflicts of interest. Affiliation, seniority, and gender are used to quantify and monitor the diversity of the community and of the TAC. The whole user portal system was implemented following the EU General Data Protection Regulation. Using their user portal identification, the user can log on to all of the ESO services, and the user portal itself acts as a gateway to these services, as shown in Fig. 3.

NSF's NOIRLab's UPP will provide publicly available information to unauthenticated users, but logging in is required to submit proposals, for time allocation process, to prepare
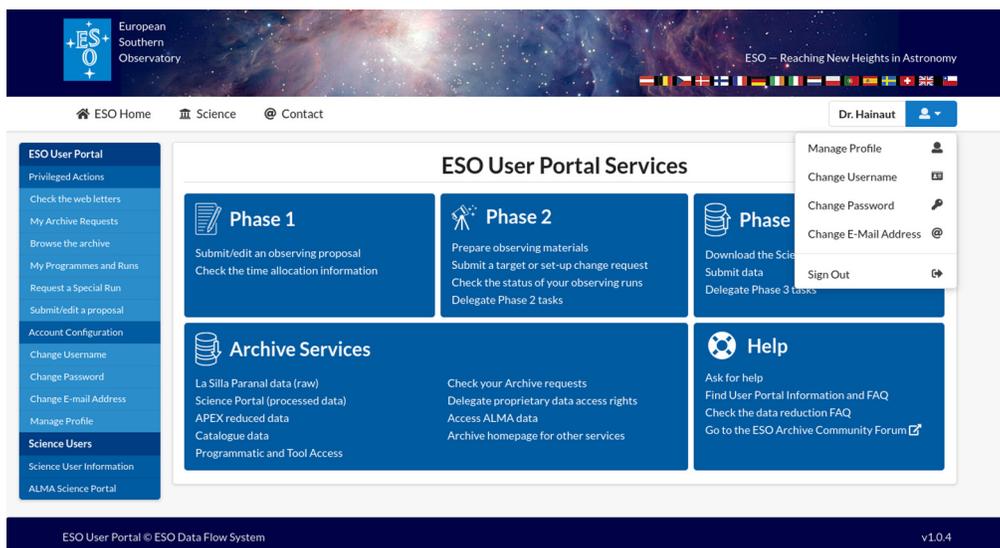

Fig. 3 The ESO user portal landing page, from which a user has access to most services provided by ESO.

observations, and to access proprietary information. The UPP will provide a mechanism to manage user roles and permissions.

Similarly, the UPP will provide a common web interface to the user-facing applications. In addition, the UPP Common Software Services are focused on providing a high-quality experience for users. There will be a user-facing dashboard along with administrative tools to configure user accounts, define collaboration groups, and authenticate data-access permissions. This system will also notify users of the status of their proposals and of approved observations that are in progress. Users will access all UPP services using a common single sign-on system. There will be a help desk system that may be used by the observatories and their partners. Over time, the help desk team will develop a knowledge base to facilitate support. Metrics will track performance, including program completion rates, archival data use, and publications.

2 End-to-End Operation Sub-Processes

In the following sections, we will discuss and compare the various tools and steps of the end-to-end operation process and how they will be implemented for the ELTs.

2.1 *Exposure Time Calculators*

A preliminary stage to any project consists of selecting the instrument (and the configuration of that instrument) that will collect data suitable to address the scientific question at hand. While experienced observers get a feeling of what magnitude is too faint to be observed, complex instruments and/or novice users require a more quantitative approach. With the exposure time calculator (ETC), the astronomer can compute the exposure time required to reach the desired signal-to-noise ratio (SNR) on their targets. The ETCs can simulate a range of objects (stars, galaxies, quasars, etc.), observing conditions (seeing, moon, and airmass), and sky transmission and brightness. They include transmission curves for all of the optical elements involved and can account for various point-spread functions (or line-spread functions for spectrographs). The ETC is meant to provide a precise estimate of the SNR for a target, whereas more complex instrument simulators can produce realistic data frames as they would be produced by an instrument, including noise features, point- or line-spread-function, etc. There is a continuum between the simplest ETC and the most complex simulator.

At ESO, the ETCs have been web-based since their early days. They are being completely overhauled to the new technology standards, and their calculation engine is upgraded with the goal for the results to reach a 10% accuracy. The new ETC web interface[16] is designed so that users unfamiliar with the instrument do not need to read the full documentation to determine whether the instrument is suitable for their needs (see Fig. 4 for an illustration). Due to the APIs
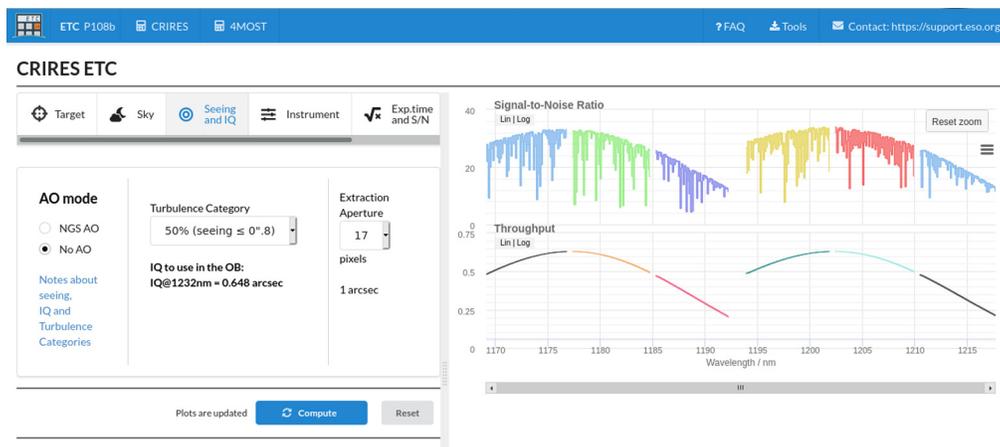

Fig. 4 An example of ETC: a zoom on a couple of spectral orders of ESO's CRyogenic high-resolution InfraRed Echelle Spectrograph (CRIRES).

and the fact that the ETC uses the same IP as the other tools, the Phase 1 and Phase 2 preparation tools (see below) will be able to call it directly to refine the SNR and exposure times. Also, the APIs allow users to programmatically scan parameters or to go through a list of targets. Finally, the APIs make the simulations available to the quality control (QC) process (see below), for instance. to compare the expected SNR with that of the observations.

NSF's NOIRLab/US-ELTP is planning to build the system to allow scientists to develop and write proposals for GMT/TMT observations based on the NSF's NOIRLab/Gemini "Explore." Potential investigators will enter only the target details and their science requirements (e.g., mode, wavelengths, and resolution) in the web-based tool, called "Prepare," which will be tightly integrated with the ETCs. They will not have to worry about which instruments they need or what GMT/TMT will have available in a specific semester. "Prepare" will automatically provide a list of configurations that will deliver the requirements that were requested, will display for each configuration how long it will take to reach the required SNR, and will present the ETC output for the selected configuration. This should help the investigators verify that they will get what they expect (see Fig. 5).
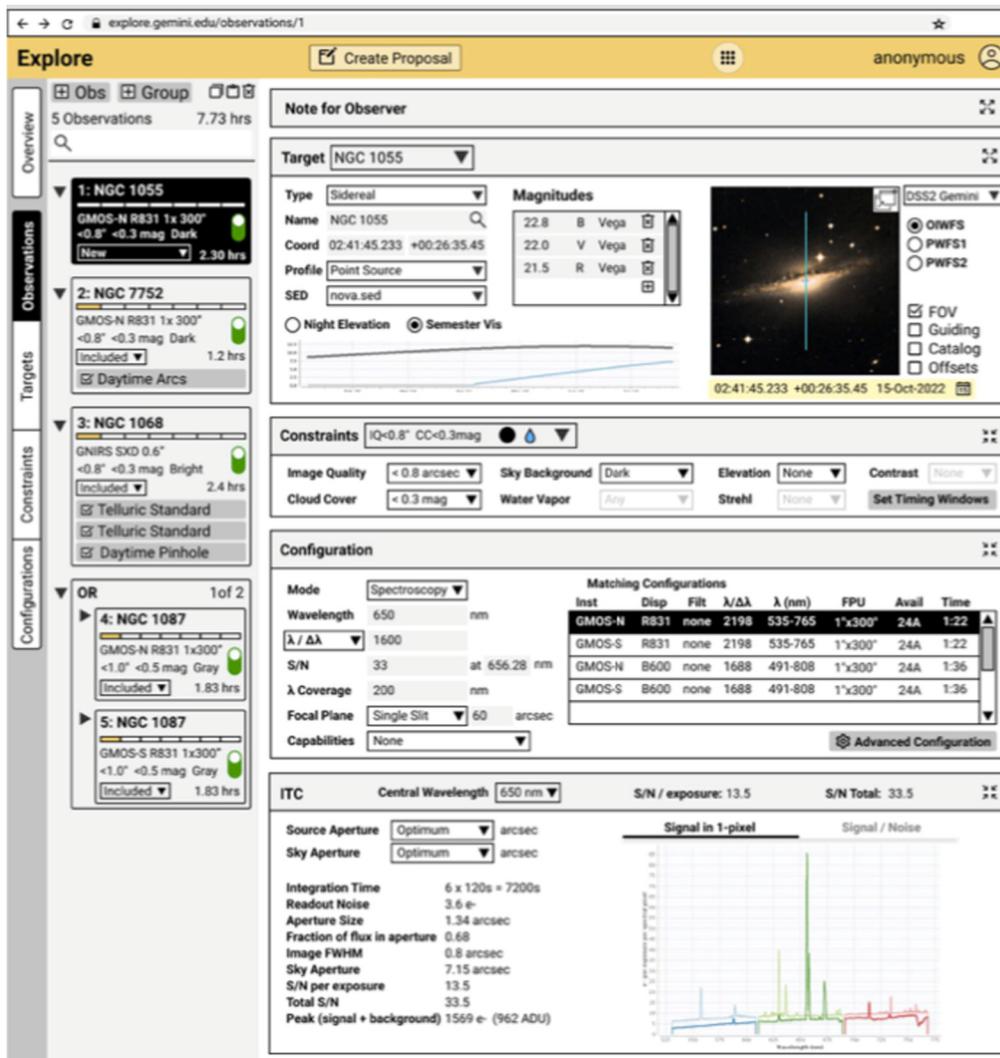

Fig. 5 Mockup of the observation view of the GPP/Explore application, the precursor of the UPP/Prepare, with a single observation highlighted and showing the list of matching configurations satisfying the high-level science requirements as well as the instrument time calculator information.

2.2 *Phase 1*

At Phase 1, the astronomer submits a proposal to the observatory, giving their science case in a narrative that will be evaluated by the scientific referees. The proposal also includes technical information on the observations that they want to perform, the instrument set-up required, and the amount of time requested.

At ESO, the former system was based on an LaTeX template with content that was parsed to feed a database. The limitations imposed by LaTeX and the maintenance issues caused by the series of organically grown tools handling the proposals triggered a complete overhaul of the Phase 1 system. The new system—P1[17]—is based on a web interface that populates the database directly using APIs. This technology choice supports collaborative editing of a proposal by all co-investigators, identified via their user portal profile. To assemble the technical information, the astronomer associates astronomical targets to instrument set-ups defined using templates from the IP (again, the same authoritative instrument model) and assigns a total exposure time and overhead. This results in simplified OBs that contain all of the information required to evaluate and schedule the observations, but without the burden of fully detailing the observations. This also ensures that inconsistent configurations simply cannot be defined. The system does the administrative work of keeping track of all of the set-ups and the total time. As all of the information is directly entered into the Phase 1 database, a variety of reports are produced; these include the traditional "proposal" document (as a PDF file) that comes in various forms, such as a fully anonymized version to be used by the referees. The web technology used in the interfaces enables various tools to be directly built into the system: visibility plots of the targets and consistency checks (see Fig. 6). Ultimately, the Phase 1 interface will be directly connected to the ETC. This will allow the user to seamlessly compute the exposure time required from their proposal and to store the results of the ETC in the proposal, simplifying the technical feasibility review by observatory staff. The Phase 1 system is designed to support all of the proposal types envisioned: the traditional 6-month cycles currently used at ESO (possibly extended to 12-month in the future), a fast-response director discretionary time program, special calls for proposals for science verification of new instruments or to cover unexpected situations, and an upcoming fast turnaround scheme with a monthly cycle.

The NSF's NOIRLab/UPP system investigators will use the "Prepare" application as described above to complete their proposals. After users enter targets and science requirements and select the instrument configurations from the list of matching options, "Prepare" will calculate the total time needed for the project using the built-in ETC and knowledge of all
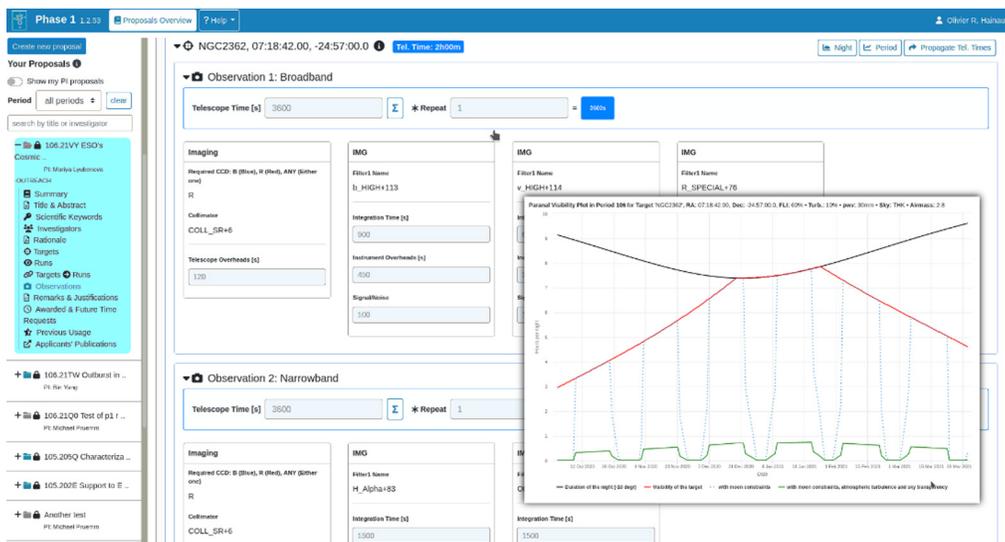

Fig. 6 Screenshot of ESO's P1 web interface, showing the definition of an observation. The inset displays the object visibility and probability of realization of the requested observing conditions over the semester.
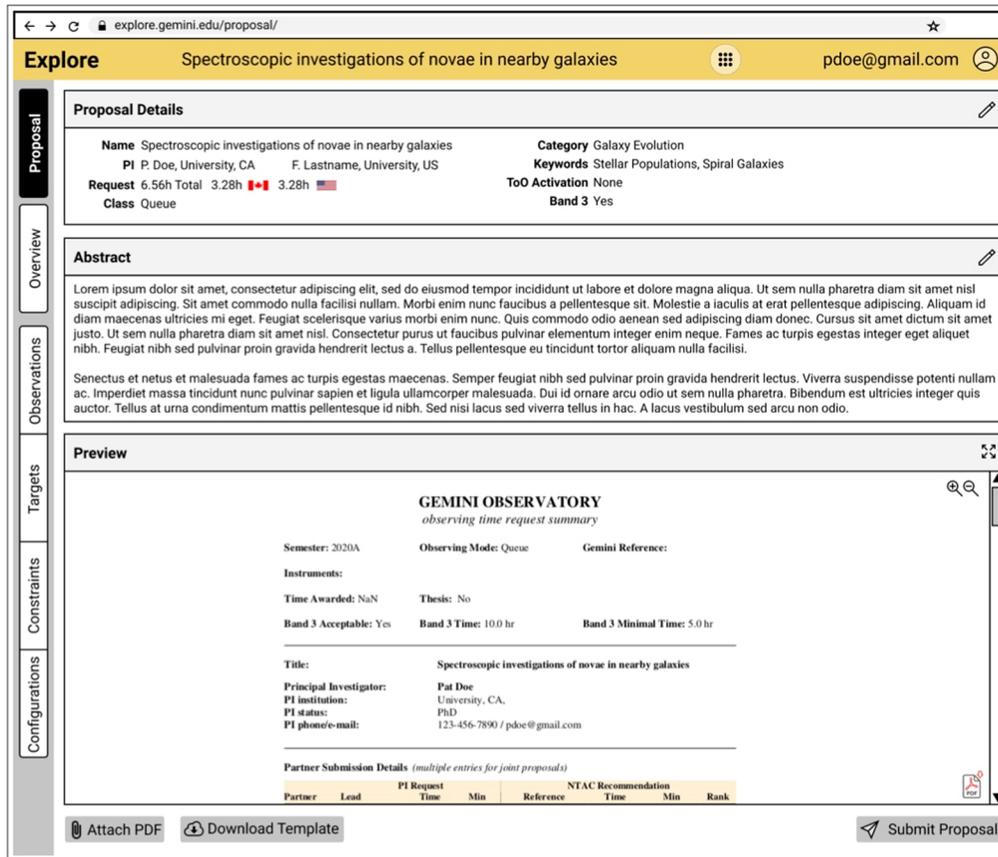
Fig. 7 Mockup of the proposal view of the GPP/Explore application, the precursor of the UPP/Prepare.

overheads. The investigators will be able to work collaboratively on the proposal, with the principal investigator (PI) being able to manage permissions. They will enter the list of co-investigators and define the time requests from each of GMTs and TMTs participants. The scientific justification, experimental design, and other "essay" sections will be created using provided LaTeX or Word templates, and the resulting PDF files will be uploaded and included with the proposal (see Fig. 7). UPP will support dual-anonymous review processes for reducing unconscious bias. The teams will be able to submit or retract proposals at any time up to the deadline. As with the ESO system, UPP/Prepare will support a variety of proposal types and multiple, simultaneous calls-for-proposals.

2.3 *Review, Time Allocation, and (Long-Term) Scheduling*

The number of proposals submitted for large telescope results in a strong oversubscription. In the case of the VLTs, the oversubscription is typically in the range of 3 to 7, and we can expect values at least as high for the forthcoming ELTs. This implies that the proposals must be thoroughly reviewed, and telescope time can be granted only to the best ones.

At ESO, the evaluation process is separated into a scientific assessment performed by an Observing Program Committee (OPC, called TAC by other organizations) composed of astronomers from the community and a technical evaluation by the observatory staff aimed at flagging proposals that are not feasible or with time requests that are wrong by a large margin. A series of new administrative tools is being developed to assemble the OPC, to keep track of suitable referees and their expertise, as well as their affiliation to flag conflicts of interest and ensure a balanced and diverse distribution of nationalities and genders. The tools will also provide the platform for the referees to access the proposals that they need to review and to deliver their grades and notes. These are then combined across the whole stack of proposals, resulting in

a consolidated ranked list of proposals that will serve as the basis for time allocation. The administration tools are currently being developed, building upon P1, which collects the proposals, and the users portal, which centralizes the personal information of proposers and referees.

At ESO, the time allocation itself is the result of scheduling the proposals starting from the highest ranked one and going down the list until all of the available time is used up, taking into account the ESO Science Operation Policies[18] and the constraints imposed by the proposal themselves (time constraints, moon, seeing, etc.). To provide additional flexibility, a moderate amount of oversubscription is allocated with a lower priority status, and programs with very loose observing constraints can also be granted time to serve as "fillers." The final schedule is reviewed by the ESO Director General, who then formally allocates the telescope time. A new time allocation system is being developed to address shortcomings of the current ones (which is slow and cumbersome and deals with each telescope independently) and to support simultaneous or coordinated scheduling of more than one telescope. This new system is designed to schedule the simplified OBs prepared at Phase 1 together with the fully detailed OBs (prepared at Phase 2, below) that are already approved (which will be critical for scheduling fast turnaround proposals).

The current operations concepts for TMT and GMT call for each partner to operate its own TAC, with inputs from the various TACs merged into a single schedule by each observatory. All of the proposals will be submitted into a common database maintained by NSF's NOIRLab, and the UPP will allow each partner to access the proposals submitted by its community of users. The partners will have the option of using the NSF's NOIRLab system to manage their TAC processes, devising their own independent process for review, or downloading the NSF's NOIRLab software and modifying it to meet their unique needs, as partners may develop their own criteria for evaluating proposals. NSF's NOIRLab's system will be designed to be flexible enough to encourage direct use of their software. NOIRLab will work with its community to develop criteria for assessing the merit of Key Science Programs and Discovery Programs, including scientific quality, data management, research inclusion, and other relevant factors. Different factors may apply to the evaluation of programs of different scales. In addition, the initial phase of review will be double-blind—proposers and reviewers will not be made known to each other.

Following their TAC process, each partner is required to provide to GMT and TMT a list of the proposals that it has approved for scheduling. This list will include the amount of time allocated and the proposal ranking. The ranked programs from multiple partner TACs must then be merged into a single integrated and prioritized observing list for each observatory. The merging TAC will include representatives from each individual TAC to assure consensus on the result. It is important to note that the multi-partner proposal merging process for GMT and for TMT will likely be different. TMT is planning a multi-partner proposal merging process similar to the NSF's NOIRLab/Gemini International TAC,[19] whereas GMT is thinking of a hybrid approach in which each partner has the option to divide their share of GMT observing time between proposals to be selected by the GMT-wide ranking subpanels and proposals to be selected and ranked through their own internal process. NSF's NOIRLab will provide the full merging and scheduling tool for TMT and a visualization tool for GMT. The observatories themselves, using these tools, will be responsible for producing an integrated observing list and assigning long-term scheduling priority to the observations.

2.4 Phase 2

The PIs who receive the good news that their proposal is approved and scheduled must then prepare the observations in full detail. Whether the observations will be performed in VM or SM simply changes the time when this is done: well in advance for SM, while more time may be provided for VM observations.

At ESO, a single web tool, P2,[20] is used for the preparation of the OBs for all visible and IR instruments on all of the telescopes. Due to the abstraction level provided by the IP, its software does not contain any instrument-specific code. In addition to the obvious advantages of maintainability and ease of adding support of new instruments, this also gives the users a homogeneous way to prepare their observations across the whole fleet of over 20 instruments. With P2, the users prepare their OBs and store them directly in the database at the Garching
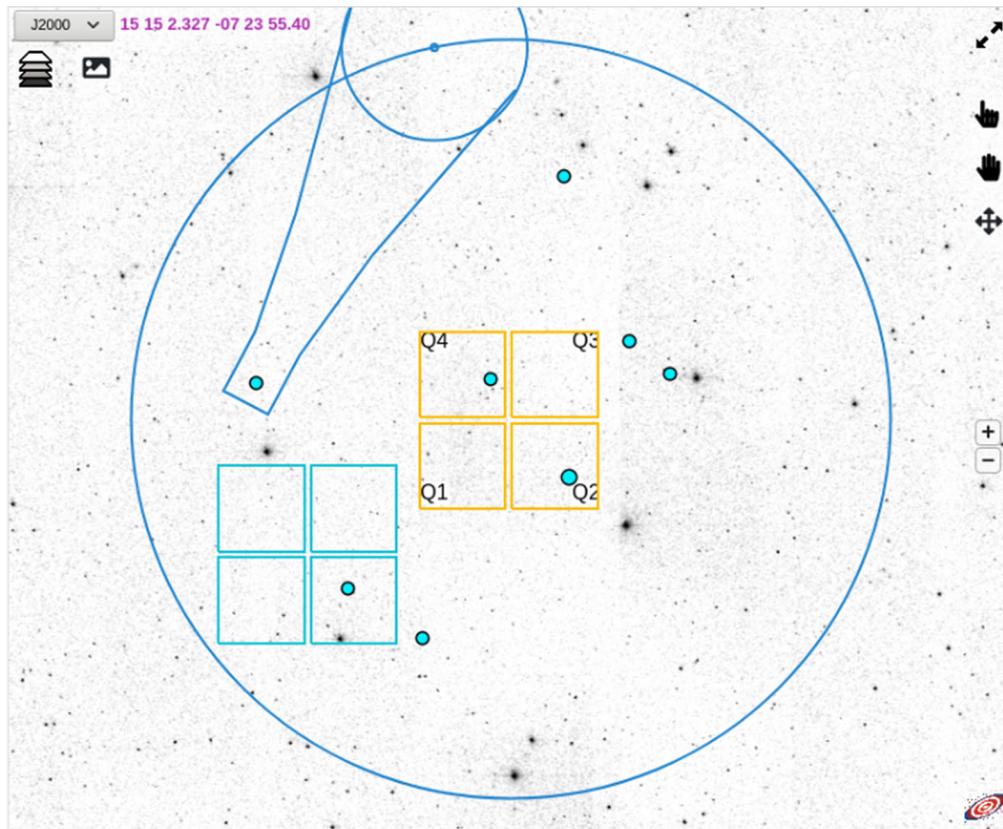
Fig. 8 A screenshot of ObsPrep, the tool used for instrument-specific observation preparation within P2 at ESO. Suitable guide stars are plotted over the digital sky survey. The user can position the guide probe on one of them, checking that it does not vignette the instrument's field of view (in orange), and select the position of an offset for measurement of the sky (in light blue).

headquarters. A bi-directional database replication ensures that the OBs are always the same in the Garching database as those at the La Silla and Paranal observatories. Furthermore, as P2 is built on documented public APIs,[21] the OBs can be prepared programmatically, either using Python scripts or custom-built user interfaces tailored to the specific needs of the observing program. Nevertheless, some complex instruments require additional preparation, with a complexity that cannot be described by the general structure of the IP templates. This preparation is done using an additional tool, ObsPrep, the interface of which is integrated within P2 (see Fig. 8). A secondary level of abstraction, which covers concepts common to several instruments (e.g., to define a complex sequence of offsets or to select a reference star) is available to all instruments in ObsPrep. Complex functions that are specific to an instrument are implemented as a "microservice" running on an ESO server and interfaced with ObsPrep via APIs. The allocation of the fibers of a multi-object spectrograph to the objects from a catalog and the calculation of the performance of an AO system based on the position of the reference stars relative to the target are examples of such microservices.

The UPP blurs the line between Phase 1 and Phase 2. Indeed, when a user prepares his/her proposal, if he/she accepts the automatically generated, default, observation in the "Prepare" application, that observation already contains all necessary Phase 2 information under the hood, including calibrations. If the default parameters are acceptable, then the user can simply change the status of the observations from "Approved" (by the TAC) to "Ready" (to be scheduled). Automatically generated sequences do not require human review, so these observations will be immediately available for the nightly scheduling process.

In the case that the full automated configuration of the observations needs to be customized, the "Prepare" allows for modification of the instrument's parameters as well as details of the observing sequence.

Fig. 9 A mockup of the observation view of the GPP/Explore application with a single observation highlighted and showing the sequence editor view.

The sequence editor panel (see Fig. 9) shows the input parameters and the resulting sequence of observations in the OB (including the acquisition and on-sky calibrations for facility instruments). Here, the user can see exactly what they will get when the observation is executed. Editing basic parameters such as spatial or wavelength dithers will cause the sequence to be automatically regenerated. Again, automatically generated sequences can be moved to "Ready" status directly without human validation. However, once a sequence has been manually edited, human validation is required before its status can be set to "Ready."

In addition, each automatically generated observation has associated calibration observations also generated automatically. The associated calibrations themselves have limited user configurability to minimize opportunities for errors. Finally, the observations may be combined into OR groups and AND groups, which may be nested. An "OR" group allows for inclusion of a larger sample of observations to facilitate scheduling. "AND" groups allow the users to request relative timing constraints between the observations in the group (see Fig. 10).

Additional information such as the science targets' descriptions and finding charts can be edited in the "target" view of the Explore application during Phase 2 (see Fig. 11). All of the features and the Phase 2 process described above will be adapted to meet GMT and TMT policies.
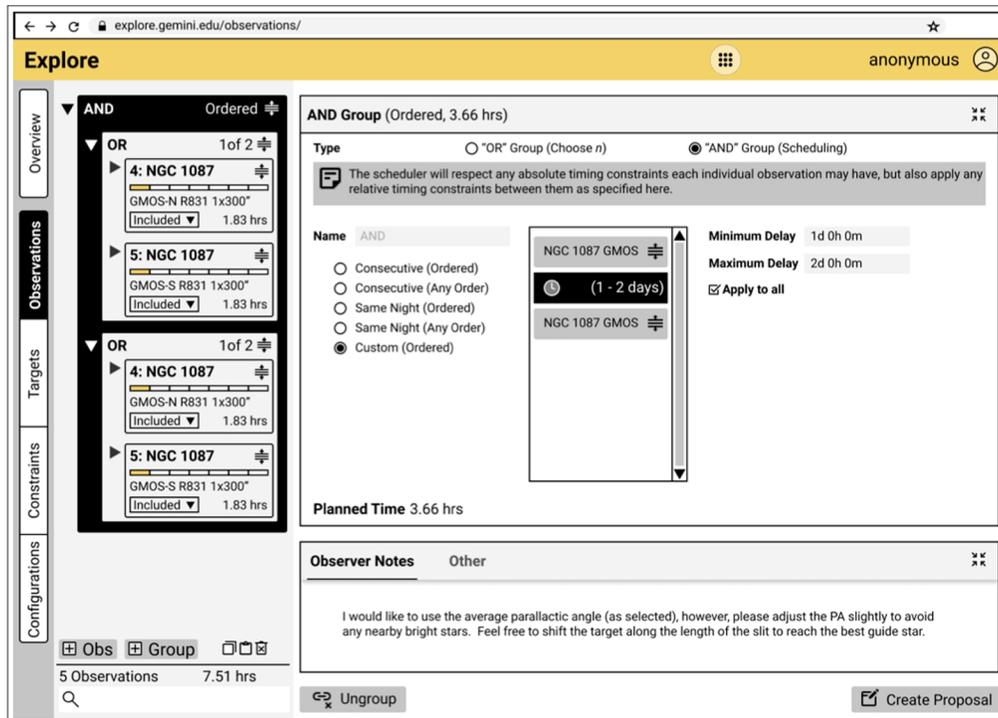
Fig. 10 A mockup of an AND group of nested OR groups in the observation view of the GPP/ Explore application.

2.5 *Short-Term Scheduling and Observation Execution*

In SM, however, the staff observer must filter and rank a large number of OBs belonging to numerous programs on various instruments, taking into account their individual visibilities, constraints (in terms of seeing or turbulence characteristics, moon illumination, humidity, etc., but also time constraints for time-critical observations), and priorities and pick the most suitable one for execution. The observatory staff performs the observations and evaluates its quality and success. It is possible to alter the planned observations (e.g., to adjust for the brightness of a flare in a ToO, or to compensate for cloud requiring a longer exposure time than originally planned), but SM observers use this capability sparingly and responsibly.

In VM, the detailed scheduling of a night of observation is performed by the visiting astronomer, possibly supported by tools provided by the observatory.

Intermediate modes in which the scientists requesting the data can interact with the observatory staff and contribute remotely to the observations are also considered. For a variety of reasons, not the least of which, for safety, the remote observer can only watch and talk, but not directly act upon the observations; this mode is therefore often called eavesdropping mode.

At ESO, in VM, this is done by the visitor populating an "execution sequence" in the P2 tool (see Fig. 12 for a screenshot) with OBs that will be executed one after the other. The execution sequence can be updated at any time; due to the real-time bidirectional database replication, this update can be done on-site, but also from any other place, enabling remote observations. Furthermore, the execution sequence can be filled and controlled using its APIs, allowing for complex strategies to be implemented via program-specific scripts. This enables a delegated visitor mode, which is essentially a VM program performed in eavesdropping mode and useful for very short and/or repetitive VM runs.

In SM, the observatory staff filters the OBs that are currently observable (in terms of visibility and observing and time constraints) and ranks them. The ranking accounts for priorities, but also for the scarcity of the observation conditions needed (so that an OB requiring, for instance, 0.5 arcsec seeing will have a much higher rank than one satisfied with 1 arcsec seeing) and of the remaining suitable observing time until the end of validity of the OB. Hence, the OB(s) rising to the top of the list (higher rank) are the most important and most urgent ones. At this time,
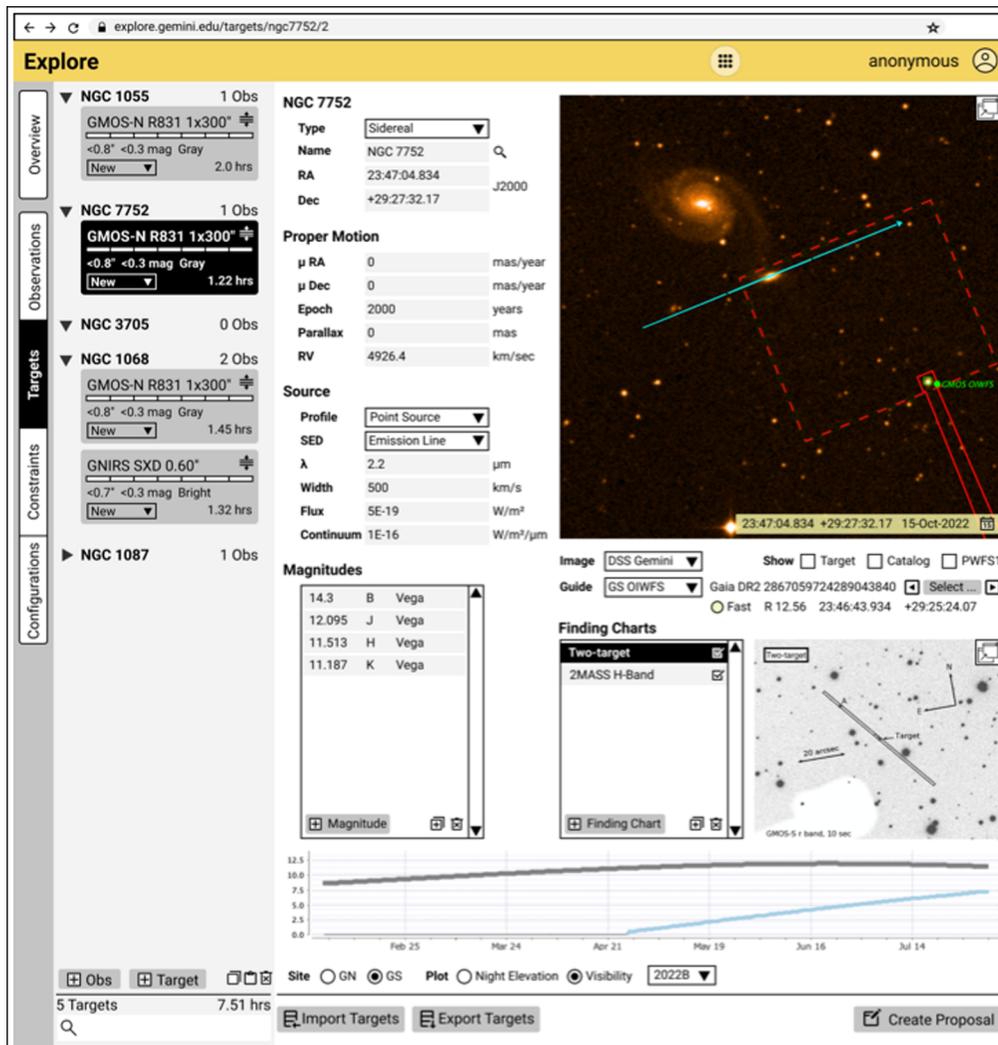
Fig. 11 A mockup of the targets view of the GPP/Explore with an observation selected.

only the current observing conditions are considered; however, ESO is planning to include short-term weather forecasts in the selection. For observations for which this is relevant, collisions with the laser beam from another telescope are also considered, accounting for the rules of engagement and priorities among the various ongoing programs. Finally, once an OB is selected for execution, it is passed to the instrument control system, where a sequencer running on the instrument workstation will execute the various steps described by the templates in the OB. The instrument returns a signal indicating that the OB is completed (or aborted, if appropriate) and sends all of the data files produced by the observations back to the DFS. This short-term scheduling is currently performed with the OT (see Fig. 13 for a screenshot), a Java desktop application. Its web-based replacement is being developed, including new ELT requirements, and will address technology obsolescence. It is scheduled for deployment in 2023.

At TMT, the observations will be supported by one support astronomer (SA), based at the sea-level control room and responsible for the instrument operations, the data acquisition, data QC, reporting, etc., and by two operation associates, operating the telescope and AO/LGS and monitoring the weather conditions from the summit. Together, this team will support the execution of science programs according to the two observing modes offered, SM and VM. In SM, the observation selection from a queue of programs is prioritized using an adaptive scheduling engine (see below). In VM, both the visiting astronomer and the SA are located at the sea-level headquarters. In eavesdropping mode, once an observation is included in the nightly plan, the users will receive a notification as well as a second notification if the observation is still in
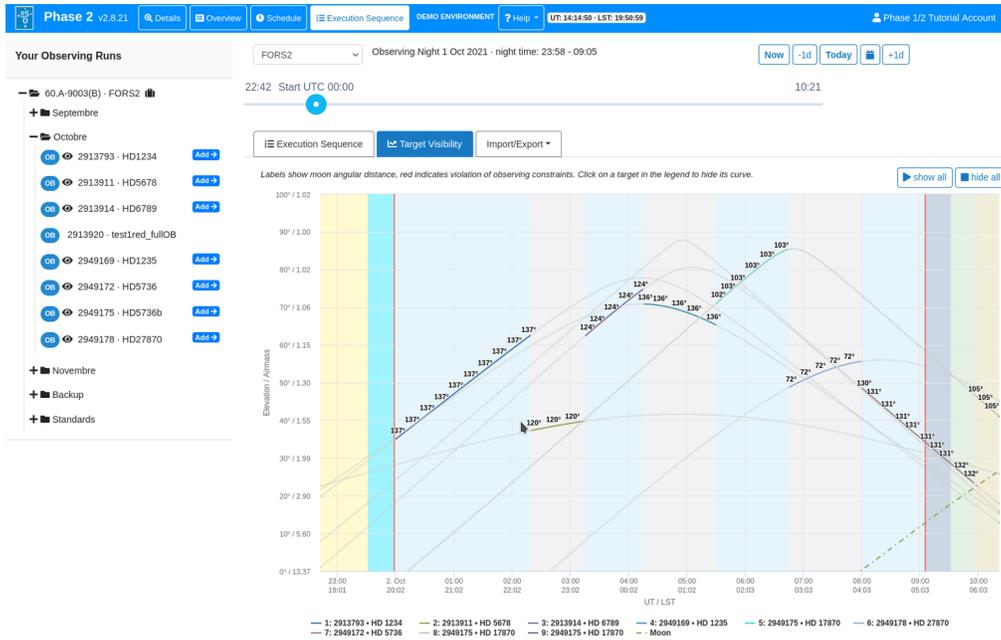
Fig. 12 Screenshot of a visitor execution sequence in ESO's P2 tool, showing the airmass of the various OBs waiting to be executed in VM.
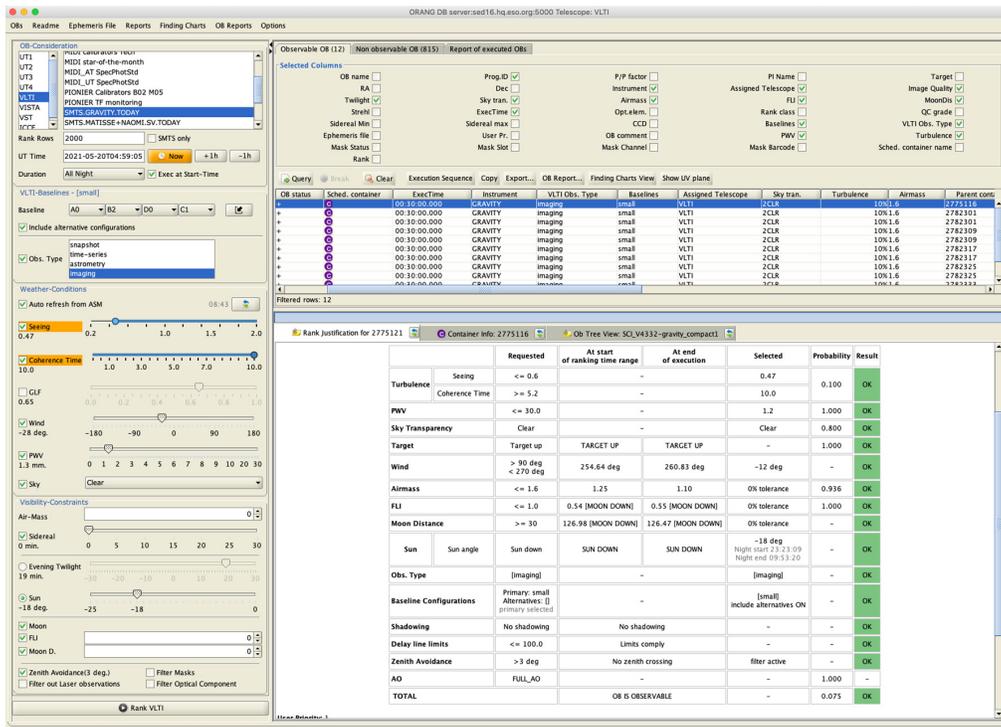

Fig. 13 A screenshot of the current VLT short-term scheduler, OT, showing a list of filtered and ranked observations. The current conditions are listed in the left column. The main window shows the ranking justification for the top observation.

nightly plan no <15 mins before the slew time. In this case, the remote users connect with the SA and assist with target identification, acquisition, or even real-time changes in observing strategy. The proportion of time used for each of these two modes will be determined by each TMT partner, based on what its user community sees as an optimal balance between VM and SM.

The adaptation scheduling engine, in addition to providing priority ranking based on current and forecast weather conditions, ranking, and completions of the programs, also balances the distribution of weather conditions and completed programs among partners, while minimizing telescope slewing time and instrument changes. The adaptive scheduling software will be provided by NSF's NOIRLab and will built upon NSF's NOIRLab/Gemini's 20 years of experience of queue.[22] One of the main requirements of this scheduling software is the automation and adaptation of SM program prioritization within minutes as events occur and as observing conditions change during the night. To help with the optimization of telescope scheduling, this software will also be able to create observing plans over multiple nights to several months in advance in simulation mode. The SA has the option to either follow the guidelines proposed by the scheduling software or bypass them if they decide that a different program should be executed instead (this decision should be documented as part of the night-report).

At GMT, the entire software framework, including the observation execution software, is highly integrated. NSF's NOIRLab's software will interface with the GMT software mostly at the level of Phase 2, where the NOIRLab-supplied Phase 2 information is converted into GMT-specific execution sequences. A queue scheduler will use Phase 2 information for pending observations; combine it with current and predicted environmental conditions, instrument availability, visibility, etc.; and prioritize the pending observations. GMT is experimenting with genetic algorithms for queue optimization. A GMTO observer will always be available to either advise the VM observing team or to perform the service observations from the queue schedule. We anticipate that VM observers will use the same queue schedule system but with a switch set to only show the team's program observations and protect proprietary information on other programs.

2.6 *Observation Log*

The historical observation logs, kept in large notebooks filled by hand, have long been superseded by digital logs filled directly using authoritative information from the instrument and telescope control systems. The detailed information about a given observation is stored directly in the FITS file, in the form of keyword-value pairs. A series of standards defining these keywords are in place, ensuring a good level of interoperability across data from different observatories. The ELTs will of course continue adhering to these standards.

At ESO, the details about the observations are collected by the night log tool and stored in a database. These include the OB details, the outcome of the online QC (QC0, see Sec. l below), and comments by the observer. Combining the information from this database with additional queries to the weather station (including measured seeing and earthquake reports if relevant) and to the problem reporting system, custom reports are generated for the VM observer, for the PIs of the programs executed in SM, for the observatory staff, etc. Alternatively, the SM program PIs can monitor the progress of their observations from a dedicated interface accessible via the user portal, as shown in Fig. 14.

For TMT, NSF's NOIRLab will provide software that will log all day and night science operations activities into a database system from which various operations- or user-oriented reports will be extracted. Fig. 15 shows a mockup of a tool, Chronicle, (designed by NSF's NOIRLab/Gemini) to extract the nighttime report. The information included in the report can be customized according to the recipient of the report (science operations staff, management, or science user). For instance, the night view of Chronicle is designed for use by the nighttime science operations team. This section starts with the usual fields for the night crew identification and other members who have provided additional support and a brief summary of the night. There is a collapsible timeline section that includes an interactive elevation plot of the executed observations, which builds as the night progresses and is color-coded by instrument. This provides an overview of all of the observations, with vertical bars indicating time lost to weather, faults, or laser shutters (due to satellites or airplanes) or spent on engineering, commissioning, or shutdown. Hovering over an observation shows a tooltip with details about the observation, and selecting an observation in the plot will highlight the timespan of the observation in the detailed section below. Double-clicking an observation will open it in the Phase 1 and 2 tool "Prepare." Below the timeline is a horizontal bar chart that breaks down the time usage for the night between
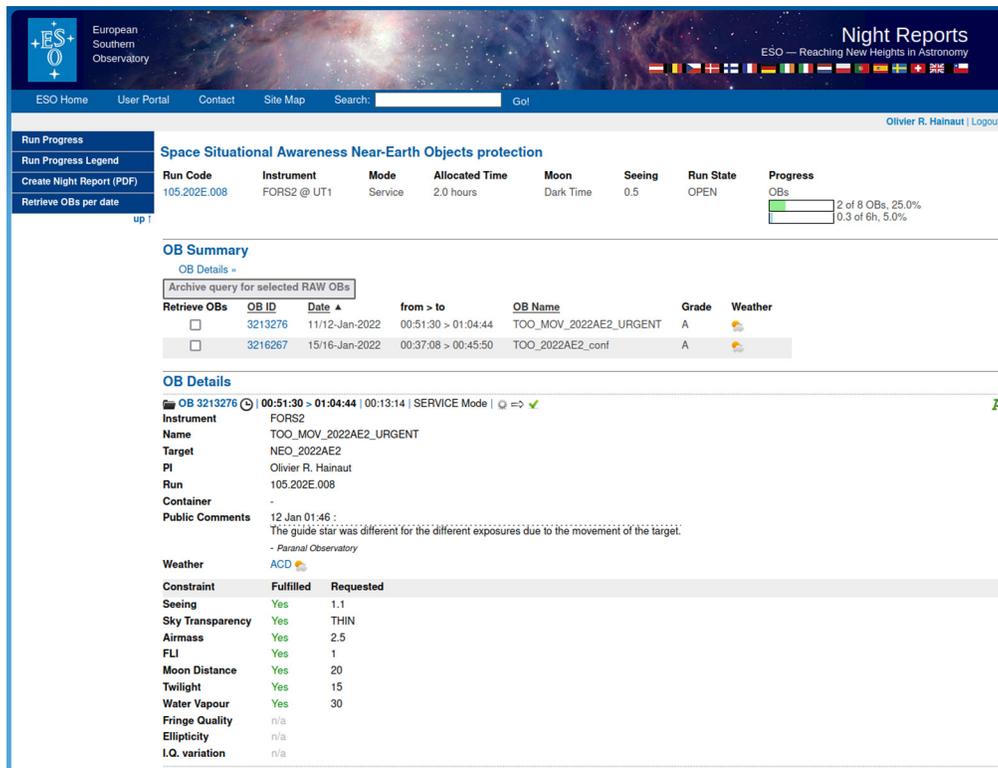
Fig. 14 Screenshot of a SM observation report produced by the ESO Night Log Tool.

science, weather, faults, engineering, laser shutters, commissioning, shutdown, and unused (idle time between programs). On laser nights. this will also display the time that the laser was shuttered for airplanes or satellites. The details section is a mix of lines added by the Night Crew and automatically populated lines with information from various sources. Each line has the time the event started, the "User" who made the entry, and a description of the event. Finally, operations metrics can be accessed through Chronicle to establish performance statistics of the science operations.

2.7 Laser Collision

Laser observations require extra precautions to protect aircraft, satellites, and even the beams of other telescopes from the laser light. The ESO ELT and GMT will use the VITRO software for aircraft to anticipate and protect aircraft. GMT and TMT will also use a transponder-based aircraft detection system to detect aircraft passing near their beams. GMT and TMT will coordinate with the laser clearinghouse (LCH) to protect against inadvertent illumination of satellites. The information from LCH will be processed by software based on that at existing LGSAO observatories such as Keck and NSF's NOIRLab/Gemini (see Fig. 16) to minimize the overheads involved in shuttering the laser due to passing satellites. Coordination with nearby observatories will adapt the Laser Traffic Control Software currently in use at other observatories, e.g., on the ESO VLT, where it is used to avoid collisions between the beams of the nearby telescopes.

2.8 Archive and Phase 3

Once the data are acquired by the instrument, they must be distributed and archived.

At ESO, a local distribution mechanism (the data handling system) takes care of making the data files available to the various systems requiring them, in particular, to the observer (for evaluation), to the online data processing pipeline (see below), and to the data transfer system, which
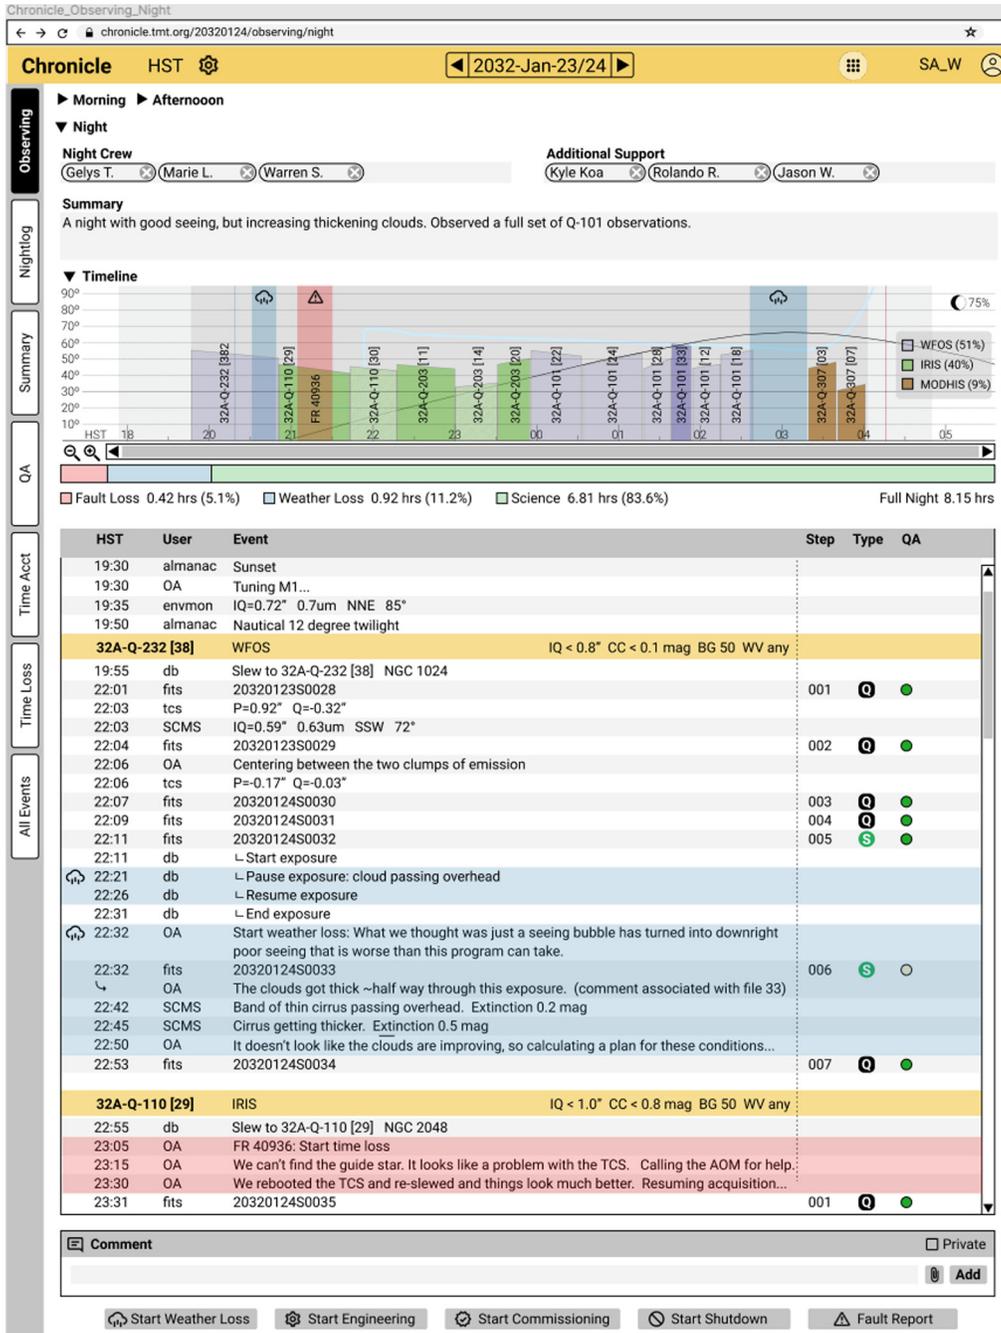
Fig. 15 TMT Nighttime Observing log mockup example based on the GPP/Chronicle.

will transfer the data to the central science archive facility (SAF) at the Garching headquarters in Germany. Although the requirement on the bandwidth between the observatory and the archive specifies that 24 h worth of data are transferred and archived within 24 h, most frames are available in the archive within 15 min. In the case of urgent observations (e.g., a target of opportunity), a mechanism can boost their priority. The daily volume of data produced by the ELT is expected to be significantly higher than that of the VLT (driven by the requirement to save a significant amount of AO telemetry data), so the data link between the observatory and the headquarters will be upgraded.

The archive holds all of the raw data and metadata collected since the early days of the DFS in the late 1990s. The storage system has been updated and upgraded several times and is likely to
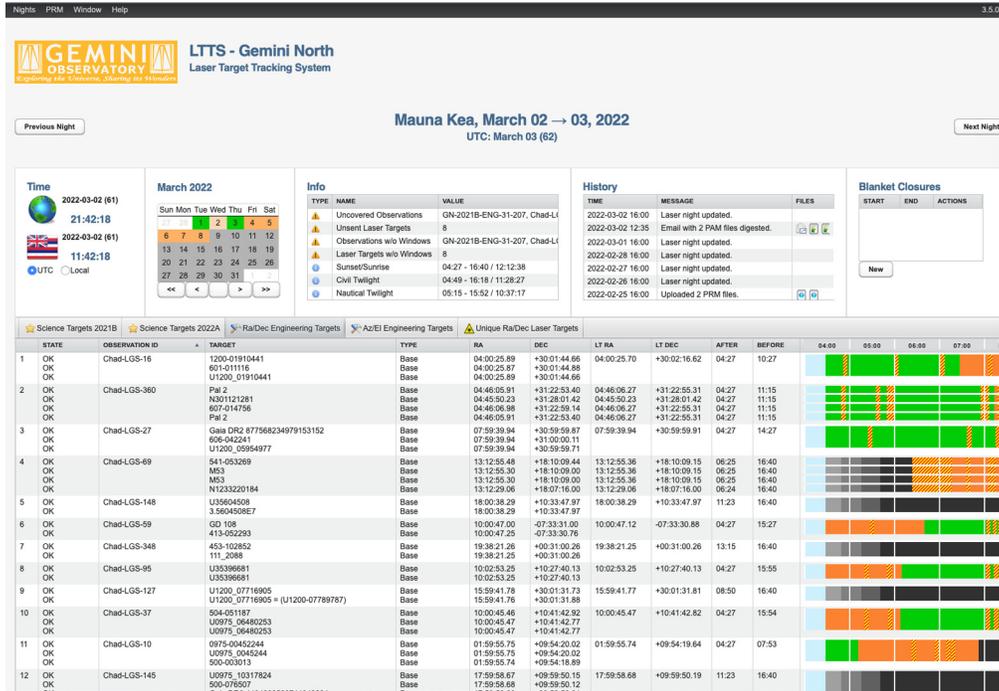
Fig. 16 NSF's NOIRLab/Gemini LTTS to automatically send target requests to LCH and automatically process the clearance windows for each target from LCH.

continue to evolve following the advances of storage technologies and possible new requirements. Originally, the archive was conceived as the data vault in which the observations would be preserved; over time, its role has been expanded to serve as the distribution center for all raw and processed data.

Indeed, the archive also contains processed data, including advanced, science-ready datasets produced either by the community (typically survey and large program data, often tailored to a specific science case) or by ESO (processed using its pipelines to remove instrument signature from the data and to produce generic science-ready products; see below) via the Phase 3 process.[7]

In addition to the traditional web forms by which users can specify their queries to the archive, a new interactive interface—the ESO science portal (see Ref. 23 and Fig. 17)—has been developed, from which the user can explore the archive using a physical description of the data among all of the instruments and telescope data archived (now also including ALMA), e.g., using wavelength and resolution, rather than filter or grating number. This was made possible due to a comprehensive data model standard adhered to for all Phase 3 processed data ingested in the archive. The science portal also shows interactive previews of both images and spectra, allowing the user to evaluate them without downloading the data. The archive can also be accessed using a programmatic interface based on the Virtual Observatory standards.

GMT and TMT science data will be transferred in near real-time to the US-ELTP science archive at NSF's NOIRLab. This will be the main access point for all science users, including archival researchers. The US-ELTP science archive will be an integrated archive for data from both GMT and TMT and will have a common interface with the other NSF's NOIRLab Data Archives. As for its other science data archive, NSF's NOIRLab will implement key International Virtual Observatory Alliance protocols for data discovery and access in the US-ELTP science archive.

GMT will retain all raw science data from instruments on the mountain for at least as long as that instrument is in service. This is to facilitate access to historical performance data even if the internet is not available from the mountain. A long-term (50 year) archive will be maintained either at the base facility in La Serena or the U.S. Science Operations Center. Redundancy and backup systems will ensure that data are not lost. Newly acquired science data should be available within an hour (goal of 5 min) of the end of the observation.
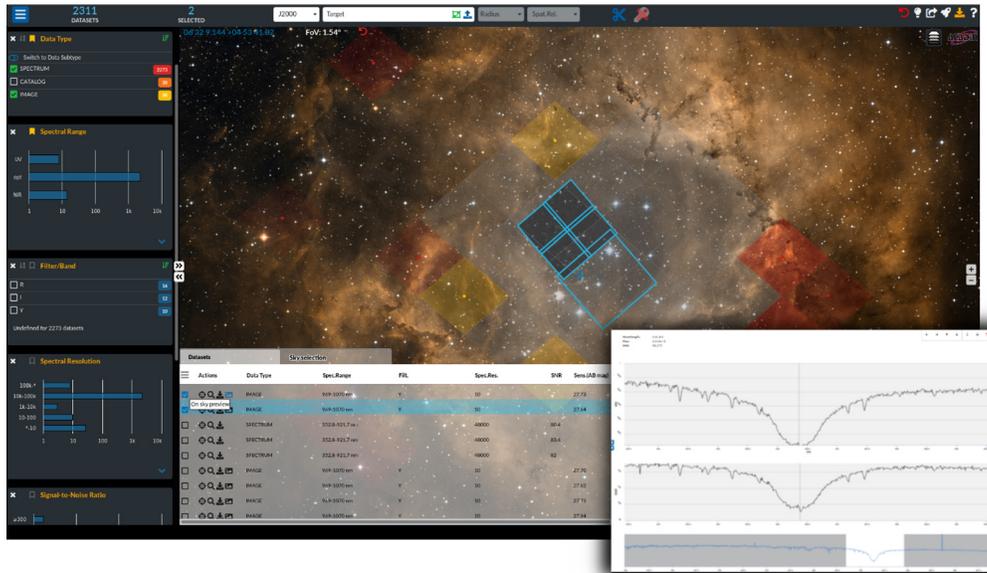
Fig. 17 A screenshot of the ESO science archive portal. Assets can be searched based on their location and on the physical properties of the data. Previews are available for most data types.

A separate engineering archive will record and curate observatory subsystem telemetry. This archive will not be copied to NOIRLab as its main purpose is to serve as a diagnostic and trending tool for troubleshooting and maintenance.

TMT uses its data management system (DMS) for handling, accessing, and archiving all TMT data over the lifespan of the observatory, which includes both science and engineering data and any external archives established by TMT partners for their own communities.

In its baseline scenario, the DMS contains two major archives: the TMT Science and Engineering Archives and the US-ELTP science archive. Immediately after their collection, the TMT raw science data will be automatically transferred to the TMT and US-ELTP archives. The calibration files associated with all instrument modes used during the night of observing will be generated the day after and will be immediately ingested into the TMT and US-ELTP archives.

In addition to the US-ELTP science archive, NSF's NOIRLab will develop a science platform, Data Reduction and Analysis Workspace (DRAW), as part of the UPP applications to serve a large and diverse community of researchers by providing a collection of tools and interfaces uniquely suited for exploring, visualizing, and analyzing data from GMT and TMT. NSF's NOIRLab will build upon its current science platforms (see Fig. 18), which have an integrated collection of high-level tools and interfaces that support discovery, exploration, visualization, and analysis of astronomical data from Rubin, Gemini, and the Middle Scale Observatories in an online environment that is co-located (either physically or virtually) with the Science Data Archive. Similarly, in the UPP/DRAW there will be pre-installed and pre-configured tools and software so that users do not need to manage their own installations. This capability will facilitate research inclusion because scientists at under-resourced institutions often have limited computing capability at their home institutions. The specific tools to be provided for working with data produced by GMT and TMT will be defined in collaboration with scientists who plan to use these telescopes and will take advantage to the greatest extent possible of the resources already provided by NSF's NOIRLab. A basic set of capabilities will likely include an online notebook environment (e.g., Jupyter or another newer technology); access to large astronomical catalogs including imaging and spectroscopic data sets through image cutout services, cross-matching, and personal allocations of file storage (e.g., Astro Data Lab provides 1TB of virtual disk storage per user); database storage (e.g., 250 GB of personal MyBD database storage in Astro Data Lab); persistent repository for published data; and computing for individual users.
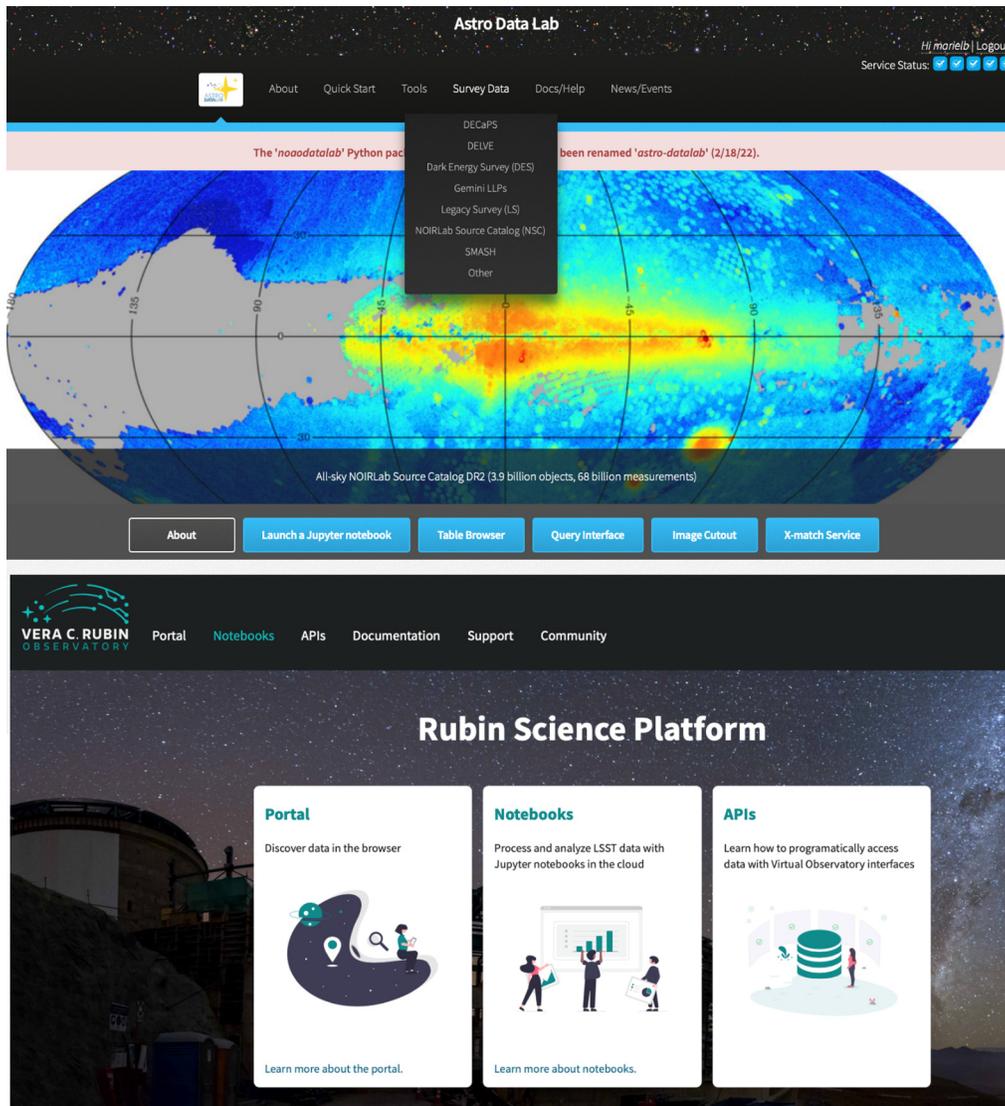
Fig. 18 Top: NSF's NOIRLab Astro Data Lab. Bottom: NSF's NOIRLab Rubin Science Platform. Both are the current precursors of the UPP/DRAW.

2.9 *Data Processing Pipelines and Infrastructure*

The complexity of the raw data produced by recent and upcoming instruments is continuously increasing. In many cases, gone are the days when a quick look at the frames with a basic exploration tool could reveal whether the observation was successful or not and when a home-made collection of scripts could process the data and allow an astronomer to start analyzing them. Data processing pipelines are now a critical component of the instrument.

At ESO, the observatory commits to providing, for every instrument, a pipeline that removes the instrument signature from the data. This includes converting the instrumental units (position in pixels, intensities in ADUs) into physical ones (pixels in sky coordinates or wavelengths, and fluxes in physical units or via a photometric zero point). This also includes re-shaping the data from whichever format they come (which can be extremely complex, for instance, in the case of integral field spectrographs or multi-object spectrographs) into one-dimensional, two-dimensional, or three-dimensional (3D) frames. At the core of the pipelines lies the Common Pipeline Library (CPL), a collection of C functions performing basic tasks, and the High-level Data Reduction Library, another series of functions to perform advanced tasks that are common to several pipelines. The pipeline for one instrument is constituted by a series of recipes, which

can be called independently (e.g., to create a master flatfield from a series of raw flats) or together via a workflow that connects them in the right sequence and selects the right input for each recipe. Currently, the pipelines are operated via various infrastructures that were developed and evolved to support the various use cases (from the telescope to the final user's computer). In preparation for the ELT, ESO is completely overhauling the infrastructure that calls the pipelines. The new infrastructure will be deployed at the telescope for quick look and QC, at the headquarters for advanced QC and production of data products, and on the user's system for final data processing. Also, the CPL is being packaged within a layer that will make it directly accessible from Python. This will allow users (both astronomers and developers) to directly call CPL functions and pipeline recipes from a Python script (with completely Python interfaces). The goal is to benefit from the power of the highly optimized CPL recipes in the rich and flexible Python environment. This will also allow developers to prototype new recipes in Python.

GMT and TMT are planning to provide seven first-generation instruments for the two telescopes combined, including both imagers and spectrometers (see "Instrument Program of the Extremely Large Class Telescopes" in this issue). Each instrument will be delivered to GMT/TMT with its set of operational documentation (user manual, calibration plans, and data-reduction manual) and data-reduction pipelines. NOIRLab is currently working with GMT/TMT to agree to common standards for data formats, metadata, and environment. The US-ELTP goals is to provide a DRP framework with three modes for each GMT/TMT instrument: (1) a quick-look mode designed for rapid turnaround and assessing data quality during observing; (2) an automated mode to run for standard data reduction (SDR), in which little or no user interaction is required; and (3) an interactive mode in which users can interact with and modify reduction parameters. The observatories (including their instrument teams) and NOIRLab will work together to define SDR-supported observing modes, SDR-required calibrations, and the procedures and data products for SDR. DRP recipes are expected to be continuously developed and improved. The evolution and improvement of the pipeline modules involved in the production of science-ready data will be the joint responsibility of GMT/TMT staff, GMT/TMT instrument team members, and NOIRLab staff involved in the development of data reduction software. The science user community also will be able to provide feedback on the DRPs. Pipeline and recipe code will be version-controlled, and changes will only be incorporated after review by all of the relevant parties.

2.10 *Quality Control and Advanced Data Products*

Traditionally, the observer controlled the quality of the data that they just acquired, adjusting their exposure time and strategy accordingly. While this is still the case in VM, the quasi-industrial aspect of SM implies a more formal approach to QC. Various levels of QC can be considered. QC0, in quasi-real-time at the telescope, as soon as the data are acquired, a first set of tests are performed, and the sequence of observations can be affected (e.g., repeating an exposure that was not up to specifications). QC1 can take place during the day and may affect the following night. Finally, QC2 is a long-term, slow-paced study that will only affect the evolution of the instrument.

At ESO, currently QC0 systematically compares the observing conditions with the specifications provided in the OB (e.g., in terms of seeing) and in some cases performs measurements on the data themselves to grade the observations (A: meet the requirements; B: almost meets the requirements and does not need to be repeated; and C: observation failed, must be repeated). QC1 focuses on the stability of the instrument within its specified performance, evaluated from the calibration frames (including parameters, such as the readout noise of the detector, the resolution of the spectrograph, the photometric zero-point of the camera, etc.). Discrepancies trigger an alert that is investigated at the observatory, with consequences such as the replacement of a calibration lamp or the adjustment of an optical element. A by-product of this QC1 step is the production of quality-certified master calibrations (flatfields and response curves) that are stored in the archive, from which they can be associated with the scientific data to which they are relevant.

In preparation for the ELT, ESO is now increasing the scope of its QC and developing a new infrastructure that will support each level of the QC process and interface with the new data processing infrastructure described above to run the pipeline. For QC0 at the telescope, the system will also interface with Scuba, the real-time QC0 platform developed at the observatory. QC1 is now taking place at the Santiago office, taking advantage of the availability of the instrument and operation specialists. QC2, at the Germany headquarters, will focus on the science QC, improvements of the pipelines, and creation of data products. With the development of science-grade pipelines, ESO has embarked on the systematic processing of the data for more and more instruments, generating science-grade data products that are made available in the archive and highly appreciated by the users.[6]

At the TMT, once the SM OB has been executed, the SA will provide a preliminary assessment (QC0) of the OB quality by comparing the ambient weather conditions with those requested by the PI. This real time QC will be complemented by measurements made directly on the image (full width at half maximum, Strehl, and SNR) and comparing them with requirements from the science proposal. As for ESO, this QC0 aims at providing triage of the executed OB between those executed successfully and those that have failed execution (due for instance to the degradation of weather conditions or instrumental problems). The day after, observatory staff will be able to provide complementary QC as needed, such as in the few cases when the QC0 assessment is flagged as uncertain.

Once the calibration files associated with the science observations are generated, they are transferred to the TMT and US-ELTP archive. This step will trigger automatic reduction of the science data using the data-reduction pipelines maintained by NOIRLab. During this process, a final check of image quality (QC1) will be provided using QC tools applied to both the calibration and science data. The result of the QC1 process applied to calibration data will be fed back to the observatory and will be used by the science operations staff to certify the calibration and monitor the status of the science instruments. The result of the QC1 applied to the science data will provide PIs with a final verification of the image quality obtained.

At the GMT, the observer will use quick-look reductions to compare the quality of the data with what was requested by the PI. This QC0 can include image quality, but also SNR and other parameters. Quick-look reductions may not use the most optimal calibrations and may skip steps in the reduction process to produce a result rapidly. Subsequent data processing using the proper calibrations and all reduction steps will provide a better QC1 assessment and may be carried out the next day. Performance parameters derived from calibration files will be automatically processed separately and compared with expectations to indicate a potential problem. Because higher level data products will mostly be produced for the US-ELTP archive by automated processes, each will have its own automated QC assessment. Manual (human) QC assessment is time-consuming and expensive and will be used only when necessary.

The US-ELTP will carry out SDR for science observations from GMT and TMT using data reduction pipelines operated at NOIRLab. The primary goal of SDR is to routinely generate reduced, calibrated data products that can be used by investigators for scientific analysis and to archive those data products for future use by archival researchers. The observatories (including their instrument teams) and NOIRLab will work together to define SDR-supported observing modes, SDR-required calibrations, and the procedures and data products for SDR.

Each TMT and GMT instrument will be responsible for defining the various data products that will be served through the archive. This will include raw data, potentially including individual reads of a detector (TMTs IRIS). It will also include data that have undergone standard data processing to remove instrument, detector, and telescope artifacts, such as bias correction, flat-fielding, etc. Some instruments, such as IRIS and GMT's GMTIFS, will produce 3D data cubes (x, y, wavelength).

Higher level data products may include extracted spectra (including individual orders from a cross-dispersed echelle or individual spectra from a multi-object spectrograph). Combined spectra (from multiple orders) or mosaicked images may be even higher-order data products.

NOIRLab will also accept contributed data sets: data that have been processed outside the UPP environment, but that provide added value, such as combined data from multiple epochs, or data requiring special processing such as precision radial velocities or high-contrast ("extreme") AO imaging.
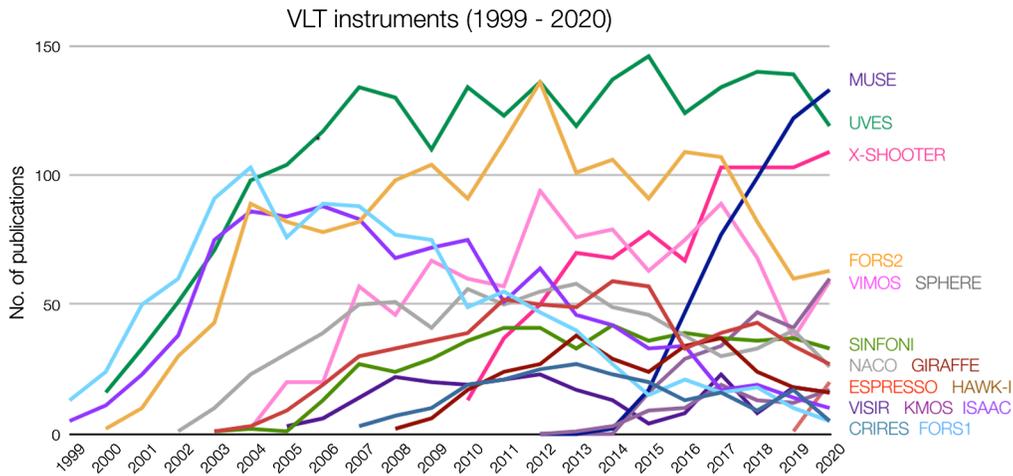
Fig. 19 Refereed publications using data from VLT instruments, using TelBib, from Ref. 25.

2.11 *Closing the Loop: Publications and Other Metrics*

The final and fundamentally crucial step of the end-to-end process is the publication of new results in scientific papers. This will foster new questions that will, in turn, trigger new proposals for observing time, closing the loop of the overall process in Fig. 1. Therefore, the number and impact of publications produced by a proposal, by an instrument, and by an observatory are important and interesting metrics to evaluate their contribution to scientific progress or more prosaically to measure the return on investment. The field of bibliometrics developed a series of analytical and statistical tools to quantify these estimates. Other aspects of the end-to-end operation process also need to be evaluated and quantified. Some can be measured from the various operation logs, such as "open shutter" duty cycle to indicate what fraction of nights assigned for science operations are spent collecting photons from the astronomical targets. For others, users must be polled, for instance to evaluate customer satisfaction and to provide specific suggestions for improvement. Finally, in each of the three ELTs, the users will also be represented in a user's committee. Each organization also collates statistics and metrics into reports to their advisory and governing bodies.

The ESO Library and Information Centre maintains the telescope bibliography database (TelBib),[24] listing the refereed papers that were produced using partly or exclusively data from ESO, obtained either through dedicated observing proposals or the science archive. The selection of the articles and the curation of their attribution to one or another facility is a complex process, described in detail in Ref. 25. This trove of data is used to put ESO's telescope in the global context of other ground- and space-based observatories and to evaluate the impact of its various instruments, as shown in Fig. 19.

NSF's NOIRLab Library Services (NLS) track refereed publications by observers using NSF's NOIRLab facilities or data products as well as refereed and non-refereed publications by NOIRLab staff. NLS maintain public libraries in Astrophysics Data System (ADS): the NOIRLab ADS bibliographic group and the NOIRLab Staff Publications public library. NLS also compile NFS's NOIRLab Publication metrics for reports (see Ref. 26). In addition, each NSF's NOIRLab facility such as NSF's NOIRLab/Gemini tracks how telescope time is used each night at each site, monitors the completion rates for observing programs, and publishes these Science Operations Statistics on the observatory website.[27] NSF's NOIRLab will track the publication of science using GMT and/or TMT data. Having a single team track metrics for both observatories will ensure consistency at least between those telescopes. Such metrics will also allow for an analysis of the scientific productivity of the different instruments on a specific telescope. They may be used to help justify retiring an old instrument to provide space for a newer, more capable instrument or to propose an upgrade to an existing instrument.

In terms of user satisfaction, ESO distributes to users and analysts a series of questionnaires and surveys, ranging from the "end of mission report" filled by visiting astronomers after their observing run to a yearly "user satisfaction survey" sent by the user support department to all
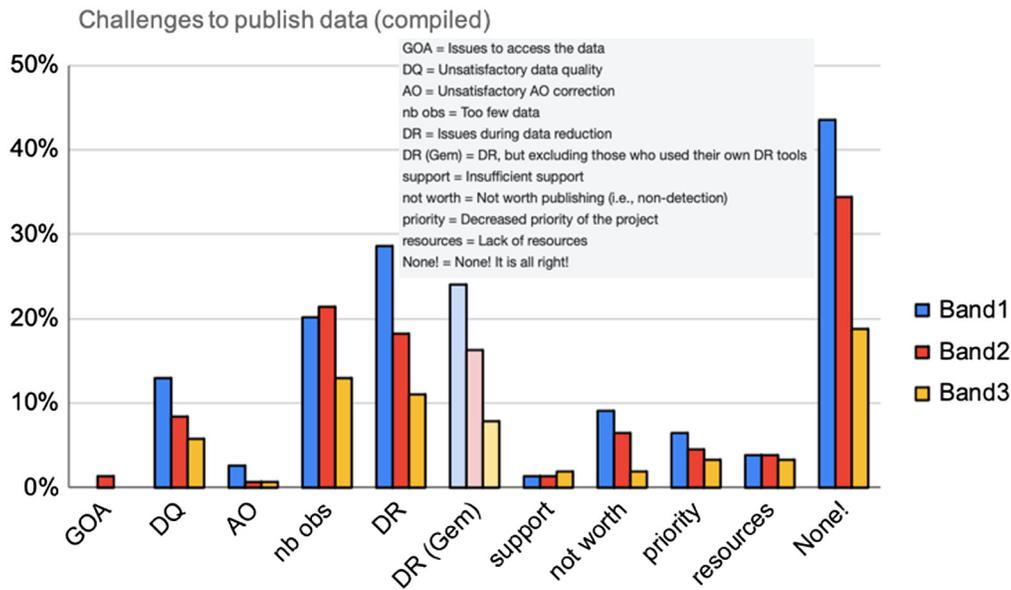
Challenges to publish data (compiled)

GOA = Issues to access the data
DQ = Unsatisfactory data quality
AO = Unsatisfactory AO correction
nb obs = Too few data
DR = Issues during data reduction
DR (Gem) = DR, but excluding those who used their own DR tools
support = Insufficient support
not worth = Not worth publishing (i.e., non-detection)
priority = Decreased priority of the project
resources = Lack of resources
None! = None! It is all right!

Fig. 20 Compilation of the responses receive between semesters 2016B and 2020B from NSF's NOIRLab/Gemini on publishing their data.

users. Ultimately, users are also represented by the user's committee, one of ESO's governing bodies. Numerous other metrics are collected and analyzed at all stages of the end-to-end process. ESO has set up a platform to collect, display, and analyze these data: the dashboard for operational metrics at ESO.

In terms of GMT/TMT user satisfaction, NSF's NOIRLab will adapt a similar approach to for NSF's NOIRLab/Gemini. Currently, NSF's NOIRLab/Gemini maintains a direct dialog with its users by sending out routine short surveys (2–3 questions) at every critical phase of user proposals/programs (Phase 1, 2, end of semester, and Phase 3). The effort has several objectives: (1) monitor the usefulness and usability of the observatory software tools and documentation, (2) determine how well the observations went, and (3) assess how satisfied the PIs are with the data and how much their expectations were met. Another objective is to identify actionable items that can improve any part of the observing process. As the name short surveys indicates, the surveys are designed to be short; they should take only a few minutes to complete. Still, for users who want to have more lengthy communications with observatory staff, the surveys always include one open question offering a text box that has no length limit. Figure 20 shows the compilation of the responses received between semesters 2016B and 2020B (total response rate of 35%) on the challenges that the users are facing that may prevent publication of their scientific data.

3 Summary and Conclusions

While the concept of end-to-end operations is ancient, its implementation as a series of interconnected subprocesses and supporting tools has been key to the success of observations at the major observatories for over 20 years. The development over the years of service and queue observing has improved observing efficiency. In parallel, observatories have been continuously reviewing their operation processes to maximize the scientific production of their equipment and increasing and improving the level of service provided to their users, reaching the level that was originally typical only for space-based observatories.

Implementing good maintenance strategies using modern tools (see the accompanying JATIS article on maintenance[28]) has helped decrease time lost to instrument and telescope faults, while helping to maintain a high level of performance. The inexorable progress of software technologies has allowed for the development of more integrated software systems, as well as the development of more user-friendly interfaces and sophisticated dataflow systems.

All three observatories have evolved toward similar solutions. New requirements are emerging from the ELTs and their instrumentation (see Ref. 29 for a description of the instrument programs of the three ELTs), from new technical capabilities and new observation strategies, and to cope with software obsolescence; thus further evolution is needed. At ESO, an ambitious overhaul of the dataflow system is taking place to support its upcoming ELT and integrate its operations with the existing telescopes. The US ELT Program also builds its end-to-end operation process on the strong legacy and experience accumulated during previous and ongoing projects.

This paper highlights the complexity of the overall end-to-end operation process for large observatories serving broad and diverse communities. The differences between the approaches deployed by the ELTs illustrate the history and legacy of various communities and organizations, but the striking similarities suggest an evolutionary convergence, based on experiments, natural selection, and of course the healthy dose of cross-inspiration to be expected among our organizations.

Acknowledgments

References

1. "Astronomical photographic plate collection," Harvard University, https://platestacks.cfa.harvard.edu/.

2. M. Mountain, "New observing modes for the next century—a summary," in *A.S.P. Conf. Ser.*, Vol. 87, pp. 235–244 (1996).

3. P. Quinn, "The ESO data management division," *The Messenger* **84**, 30–33 (1996).

4. P. Puxley et al., "The Gemini Observatory science operations plan," *Proc. SPIE* **3349**, 63–74 (1998).

5. F. Primas et al., "Fifteen years of service mode operations: closing the loop with the community," *The Messenger*, **158**, 8–15 (2014).

6. M. Romaniello et al., "The growth of the user community of the La Silla Paranal observatory science archive," *The Messenger*, **163**, 5–9 (2016).

7. M. Arnaboldi et al., "Phase 3 status and access to science data products from ESO public surveys," *The Messenger*, **156**, 24–29 (2014).

8. M. F. Sterzik et al., "The scientific return of VLT programmes," *The Messenger*, **162**, 2–8 (2015).

9. F. Patat et al., "Period 100: the past, present and future of ESO Observing programmes," *The Messenger*, **170**, 5–10 (2017).

10. O. R. Hainaut et al., "End-to-end operations in the ELT era," *The Messenger*, **171**, 8–13 (2018).

11. P. Hirst and R. Cardenes, "The new Gemini Observatory archive: a fast and low-cost observatory data archive running in the cloud," *Proc. SPIE* **9913**, 99131E (2016).

12. K. Labrie et al., "DRAGONS - data reduction for astronomy from Gemini Observatory north and south," in *ADASS XXVII, ASP Conf. Ser.*, J. Ibsen et al., Eds., Vol. 522, p. 321 (2019).

13. S. Margheim et al., "Past and future evolution of Gemini operations," *Proc. SPIE* **10704**, 107041P (2018).

14. B. Miller et al., "Gemini OCS upgrades: a blueprint for NOIRLab operations software in the 2020s," *The Mirror*, **2**, 42–44 (2021).

15. "ESO user portal," https://www.eso.org/UserPortal/.

16. "ESO exposure time calculator," https://etc.eso.org/observing/etc/crires

17. "ESO phase 1 tool," https://www.eso.org/p1demo.

18. "ESO optical/infrared telescopes science operations policies," ESO Council, Cou-1947, 2020, https://www.eso.org/public/about-eso/committees/cou/cou-154th/external/Cou_1847_rev_Science_Policies_050520.pdf (2020).

19. M. Lemoine-Busserolle et al., "The Gemini Observatory International Time Allocation Committee (ITAC) process," *Proc. SPIE* **11449**, 114491A (2020).

20. "ESO phase 2 tool," https://www.eso.org/p2demo.

21. "ESO phase 2 API," https://www.eso.org/sci/observing/phase2/p2intro/Phase2API.html.

22. B. W. Miller and R. Norris, "Gemini queue planning," *Proc. SPIE* **7016**, 70160V (2008).

23. "ESO science portal," https://archive.eso.org/scienceportal.

24. "Telescope bibliography database," ESO TelBib, https://telbib.eso.org.

25. U. Grothkopf et al., "Basic ESO publication statistics," ESO report (2021).

26. NOIRLab Publication Metrics, https://noirlab.edu/science/library/publications/metrics.

27. S. Leggett and A. J. Adamson, "Gemini Observatory: science operations statistics," https://www.gemini.edu/observing/science-operations-statistics (2022).

28. M. A. Ortega et al., "Maintenance concepts for ELTs," *J. Astron. Telesc. Instrum. Syst.* **8**(2), 021506 (2022).

29. W. Skidmore et al., "Instrument programs of the extremely large class telescopes," *J. Astron. Telesc. Instrum. Syst.* **8**(2), 021510 (2022).

Olivier R. Hainaut is an astronomer at ESO. He has been responsible for operations of the NTT on La Silla, then of the La Silla Observatory, and then of the VLT. He is now an end-to-end operations scientist for ESO. He received his PhD in physics from the University of Liege, Belgium, in 1994.

Marie Lemoine-Busserolle is a system scientist of the US ELT Program at NOIRLab. She received her PhD in astrophysics from the Université de Toulouse, France, in 2003.

Christophe Dumas is an observatory scientist and head of operations for the Thirty Meter Telescope. He received his PhD in astrophysics from the Université de Paris, France, in 1997.

Robert W. Goodrich is an observatory scientist for the GMT. He received his PhD in astronomy and astrophysics from the University of California at Santa Cruz in 1988. He has worked at Keck, Palomar, McDonald, and Lick Observatories as well as STScI.

Bryan W. Miller is an astronomer and lead scientist for Science Operations Development for Gemini Observatory at NOIRLab. He received his PhD in astrophysics from the University of Washington in 1994.

Michael F. Sterzik is an astronomer and head of the Data Management and Operations Division at ESO. He received his PhD in astrophysics from the University of Tübingen, Germany, in 1993.

Thomas Bierwirth is the software project manager and architect of the dataflow system at ESO. He received is master's degree in physics from the University of Osnabrueck, Germany, in 1994.

Sidney Wolff is an astronomer emeritus at NSF's NOIRLab. She led the design and development phases of NOIRLab's portion of the US-ELTP. She received her PhD in astronomy from UC Berkeley in 1966.

Warren Skidmore is the Instrumentation System Scientist for the Thirty Meter Telescope. He received his PhD in astrophysics from Keele University, UK, in 1998.

Kim Gillies is the software architect for Thirty Meter Telescope. He has also worked in operations and software at STScI, Gemini, and Kitt Peak observatories. He has an MS degree in astronomy from University of Iowa and MS degree in computer science from University of Arizona.

Biographies of the other authors are not available.